WARP3d, a Three-Dimensional PIC Code for High-Current Ion-Beam Propagation Developed for Heavy-Ion Fusion

D. P. Grote *I. Haber*
A. Friedman *S. M. Lund*

Introduction

Heavy-ion induction accelerators are being developed in the United States as the principal candidate "drivers" for inertial confinement fusion (ICF) power production. They also represent an attractive option for a high-yield microfusion research facility. Heavy-ion drivers are attractive in both cases because they promise high reliability, high repetition rate (several pulses per second are required), greater efficiency (up to 30%), lifetime capacity (a power plant would need to last 30 years, or 10^{10} pulses), and sufficient cost savings. With indirectly driven targets, heavy-ion drivers allow a favorable fusion chamber geometry; furthermore, the final focus is accomplished with magnetic fields that are immune to damage by target debris, in contrast to the final optics for a laser.[1]

The requirements for high-gain ignition of the fusion targets control the design of the accelerator. A major constraint on the beam is the requirement of a small focal spot size on the target with a large distance between the final focus lens and the target, putting a tight limit on the temperature of the beam. A thorough understanding of the beam behavior as it propagates through the accelerator is necessary to maintain the required low temperatures.

Since the space-charge–dominated beams in an induction accelerator are effectively non-neutral plasmas, analysis of the beams can be performed with computational modeling techniques related to those used in both the accelerator and plasma physics fields. Because the beam resides in the accelerator for relatively few plasma oscillation periods, particle-in-cell (PIC) simulation techniques are especially effective and have proved valuable in the design and analysis of both ongoing experiments and future machines and in the study of basic physics issues. The three-dimensional (3-D) electrostatic PIC simulation code WARP3d[2,3] was developed to study beam behavior from first principles in realistic geometries. This article presents an overview of the WARP3d code, describing in detail some of its major features and present applications in a number of areas, and highlighting the important results.

Overview of WARP3d

WARP3d was designed and optimized for heavy-ion fusion (HIF) accelerator physics studies. It combines the PIC technique commonly used for plasma modeling with a description of the "lattice" of accelerator elements. In addition to its 3-D model, a simple beam envelope, or transverse root-mean-square (rms) size, analytic equation solver was included, primarily to obtain an initial beam well matched to the accelerator lattice. The code uses the BASIS development and run-time system, which affords a flexible and powerful interpretive user interface.

The ion beam is represented by a small number (relative to the number of beam ions in the physical system) of interacting particles whose motion is controlled by the Newton–Lorentz equations of motion. The fields affecting the particles include both the self-fields of the beam and the external fields from the lattice. The self-fields are assumed to be electrostatic and are calculated in the beam frame via Poisson's equation on a Cartesian grid that moves with the beam. The electrostatic assumption is valid since, with the high mass of the beam ions, the beam is nonrelativistic throughout the accelerator, its velocity only reaching roughly 10% of the speed of light at the end of the accelerator.

There are two "field solvers" to calculate the self-fields: (1) The fastest field solver, which is generally used for long-time simulations, uses Fourier decomposition with fast Fourier transforms (FFTs). (2) An iterative solver uses successive overrelaxation (SOR) with even-odd ordering[4] and allows inclusion of complex electrode geometry. For the FFT field solver, the 3-D transform is sine-sine-periodic, the transverse boundary is a rectangular

metal pipe, and the longitudinal boundary is periodic. Capacity matrices[5] can be used to include simple electrodes in the FFT field solver to handle round pipes, structures independent of the axial coordinate, and isolated internal structures. Another variant uses a tridiagonal matrix solver along the axial direction. This method is slightly faster, but at present is restricted to conducting boundaries on the longitudinal ends of the grid.

The SOR field solver, although significantly slower than the FFT field solver, allows simulation of a much broader class of problems with a more realistic model of the conductors. This field solver was implemented using a technique, described below, to obtain subgrid-scale resolution of the conductor boundaries. With both field solvers, transverse symmetry can be exploited, simulating only half or one quarter of the system.

The external fields are applied via a lattice, which is a general set of finite-length accelerator elements, including quadrupoles for focusing, dipoles for bending, induction gaps for accelerating, elements with arbitrary transverse multipole components, and curved accelerator sections (see Ref. 2). The transverse multipole components specify the applied fields, in one form, as an electrostatic or magnetostatic potential represented by an infinite sum of terms expressed in cylindrical coordinates with the direction of beam propagation along the longitudinal axis, z.

$$\varphi = \sum_{n,\nu\geq 0} V_{n,\nu}(z) r^{n+2\nu} \cos(n\theta + \nu) \qquad (1)$$

The $V_{n,\nu}(z)$ determine the contribution of each term. The electric or magnetic fields are calculated by taking the appropriate derivatives. As an example, a quadrupole field is given by the term with $n = 2$, and $\nu = 0$. Terms with $n > 2$ or $\nu > 0$ are referred to as high-order terms.

The field from the lattice elements can be specified at one of several levels of detail—from hard-edged, axially uniform fields at the lowest level, to self-consistent inclusion of the conducting electrodes of electric elements in the field solution at the highest level.

Aspects of WARP

A number of features and capabilities have been implemented in WARP3d to enable simulation of a broader class of accelerators, both by way of inclusion of the relevant physics and by optimization of the code to allow simulation of larger-scale problems. This article describes the most important and novel of these in detail.

Lattice and Fields

A hard-edged field is the lowest level of description of lattice element fields—one that is uniform axially within the extent of the element and that has a sharp cutoff at the ends of the element. As the particles enter and exit hard-edged, axially uniform elements, the force applied on the particles is scaled by the fraction of the timestep spent inside the element. These "residence corrections" are used so that the particles receive the correct impulse from the element and allow larger timesteps to be used. Without the residence corrections, the forces applied to the particles would depend on the number of timesteps they spend inside the element; keeping the resulting error small would require a small timestep.

At a higher level of detail, fields that vary axially can be applied. This allows application of realistic fields, including fringe fields and the higher-order multipole moments of a magnetic quadrupole element, for example. In both the axially uniform and nonuniform descriptions, the applied field can be an arbitrary sum of multipole components, as given by Eq. (1).

The highest level of detail is specification of the field via a 3-D grid to allow application of fields that are not amenable to a multipole description, e.g., a field with many nonlinear components. The field can be calculated outside of WARP3d; for example, the magnetic field from a complex magnet design can be obtained from TOSCA[6] and input into WARP3d. With electric elements, the field can be calculated self-consistently with WARP3d by inclusion of the electrodes in the field solution, thereby obtaining the realistic fields and image effects from the electrode on the beam. This highest level of description allows modeling of complex electrode geometry and electrodes with time-varying voltages.

When electrodes are included in the field solution, to avoid the problems with the jaggedness of representing tilted or curved electrodes on a Cartesian grid, the SOR field solver allows subgrid-scale resolution of the conducting boundary locations as shown in Fig. 1. For grid points just outside the surface of a conductor, the form of Poisson's equation is changed to explicitly include the location of the surface. Normally, the evaluation of Poisson's equation at these points would require the potential at grid points that are just inside the surface of the electrode. The values inside are replaced by a value that is extrapolated from the potential at the surface and at the grid point just outside the surface. The method is similar to the method used in EGUN,[7] extended to three dimensions and time dependence.

Subgrid-scale resolution of the conducting boundaries allows more accurate modeling of electrodes for a given grid cell size, especially ones with curved surfaces, and it does so while allowing larger grid cell sizes. As an example, the fields from an electric quadrupole are produced from four parallel cylinders with appropriate voltages. Without the subgrid-scale resolution of the cylinder surfaces, the quadrupole focusing strength is reduced significantly since the jagged surface, produced by representing the cylinder by grid points only within

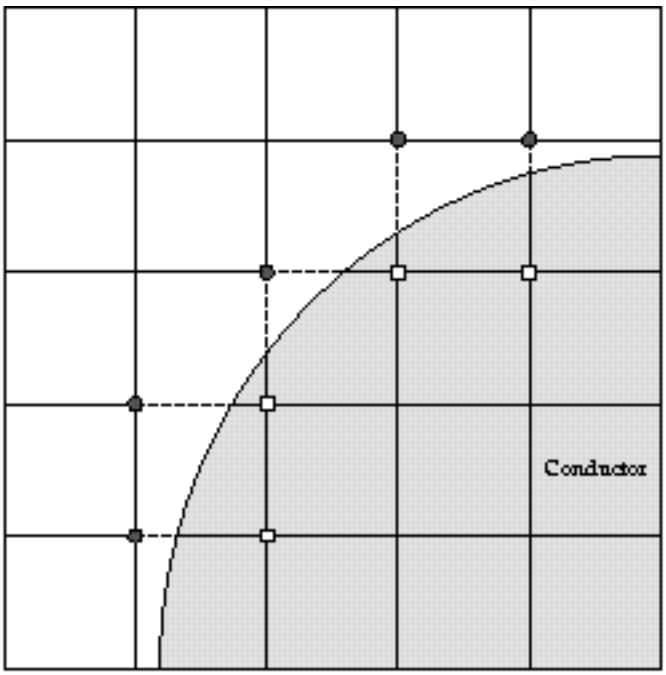

FIGURE 1. An example conductor on a grid. The numeric form of Poisson's equation is altered for the grid points at the dark circles. The equation would normally require the potentials at the nearby grid points, which are inside the conductor, marked by squares. The potential inside is replaced by a value extrapolated from the potential at the point outside and at the surface. (08-10-0896-1906pb01)

the cylinder, is effectively farther from the axis than that of the true surface. The reduced focusing changes the beam envelope on the order of 10%. A much finer grid would be needed to produce the same accurate focusing strength as that produced with the subgrid-scale resolution.

Beam Injection

Several models of particle beam injection are used, all of which can model emission from planar or curved surfaces. The first model injects particles from a surface at a constant current with a specified energy, modeling a beam emerging from a wire mesh or from a plane of uniform potential. It can be used to model injection from an emitting surface if the specified energy is low compared with the diode energy.

Two self-consistent models use space-charge–limited injection and are intended only for injection directly from emission surfaces. The beam current that is injected is not predetermined, but is calculated self-consistently from the conditions near the emitting surface to give space-charge–limited injection. The first model uses the Child–Langmuir[8] relation in a region in front of the emitting surface to calculate the emitted current. The length of the region is small compared with the length of the diode. The Child–Langmuir relation gives the space-charge–limited current in a 1-D diode. In the second model, the surface charge required to produce the desired zero-normal electric field on the surface is calculated assuming an infinite conducting plane. The amount of charge injected into a grid cell is then the surface charge integrated over the grid cell minus the charge already in the grid cell from earlier times.

With space-charge–limited injection, WARP3d can be run in an efficient iterative manner that captures only the steady-state behavior. The algorithm is similar to that used in the 2-D electron gun design code EGUN.[7] At each iteration, a group of particles is tracked through the diode. Along the paths of the particles, charge is deposited on a grid to build up a representation of the charge density of a full beam. The fields are recalculated, including the new charge density, giving the new fields through which the next group of particles is tracked on the next iteration. On the initial iteration, the particles are tracked through the fields given by the solution of Laplace's equation in the diode geometry. The process converges to the correct steady-state flow in typically 10 to 20 iterations. This capability was added to the code entirely through the BASIS interpreter and required no modification of the existing code.

Large Time-Step Algorithms

Under conditions that arise commonly in WARP3d applications, the applied focusing, bending, and accelerating fields vary rapidly with axial position, while the beam's self-fields (which are comparable in strength to the applied fields) vary smoothly. In such cases, it is desirable to employ timesteps that advance the particles over distances greater than the characteristic scales over which the applied fields vary. An example technique, used in WARP3d since the code's inception, is the residence correction described earlier. Research at Lawrence Livermore National Laboratory (LLNL) is extending this concept to cases where the effects of extended fringe fields and other smooth but rapid variations must be accurately captured, necessitating very small timesteps when a conventional "leapfrog" advance is used.[9]

Several related techniques are being evaluated: subcycling of the particle advance relative to the field solution (i.e., taking N substeps between each major timestep on which the self-consistent field is computed and applied, where N is a small integer); use of a higher-order time-advance algorithm; and force-averaging over the velocity-advance step by integration along approximate orbits (a generalization of the residence correction technique). The most straightforward of these techniques is particle subcycling. We are encouraged by subcycling tests that advance a particle through the extended field's quadrupole component of a pair of permanent-magnet quadrupole lenses, and we are in the process of implementing this technique in WARP3d.

Since the field solution and particle advance often use comparable amounts of computer time, subcycling with $N = 4$ will result in roughly a twofold speedup of the code. For some problems, the field solution dominates the run time and larger speedups can be expected.

Thick-Slice Model

Several simulations have identified the existence of short longitudinal-wavelength oscillations that can occur on a long beam. These short-length phenomena have been found to be particularly significant when the beam transverse temperature is somewhat greater than the longitudinal temperature.[10–13] In this case, the temperature anisotropy is a source of free energy that can drive the short wavelength modes unstable, leading to an "equilibration" of the temperatures. Short longitudinal-wavelength oscillations have also been seen (in simulations) in a recirculator with external nonlinearities (see the section entitled "Recirculating Induction Accelerator"). To investigate these phenomena numerically, therefore, it is important to examine the physics of a space-charge–dominated beam with numerical resolution in the longitudinal direction sufficiently fine to resolve phenomena that occur on the scale of the beam diameter. This requirement can be computationally expensive to satisfy if the beam is to be tracked through a realistic length of accelerator structure.

A new numerical technique currently being developed facilitates the investigation of such short longitudinal-scale phenomena. This approximation, which we call the "thick-slice" model, exploits the separation of scales that occur because of the locality of the phenomena being studied relative to the global beam variation along the accelerator structure. We assume that all the particles across a finite slice of the beam simultaneously experience the same external forces, in a manner similar to what is assumed in a single-slice (or 2-D transverse) model. However, the self-consistent beam dynamics within the slice are followed, including the longitudinal variation of the fields. The longitudinal particle and field boundaries at the ends of the slice are further assumed to be periodic.

Parallel Processing

WARP3d was originally designed to run on vector supercomputers and on workstations. With the present trend toward using massively parallel processing, WARP3d was ported to a parallel environment. Currently, WARP3d can run on the CRAY T3D, using both the PVM[14] message passing library and the native shared memory library, and on multi-workstation clusters. In the conversion, the interactive nature of the code (through the BASIS interpreter) was maintained. The user runs the code on a local serial workstation and interacts with it through the interpreter. That code then creates processes on the parallel machine to do the computational work. Via message passing, the user controls the progress of the simulation on the parallel machine and inquires or sets any of the code variables in a manner similar to that used when the code runs on a serial machine.

Using a model problem with up to 32 processors, timings were made by following a beam through five hard-edged lattice periods. The speedup of the code with an increasing number of processors was nearly linear. Figure 2 shows timings of the field solver and particle mover. The field solver has a slightly superlinear speedup, which is likely due to better cache memory use with the data spread out in smaller pieces on more processors.

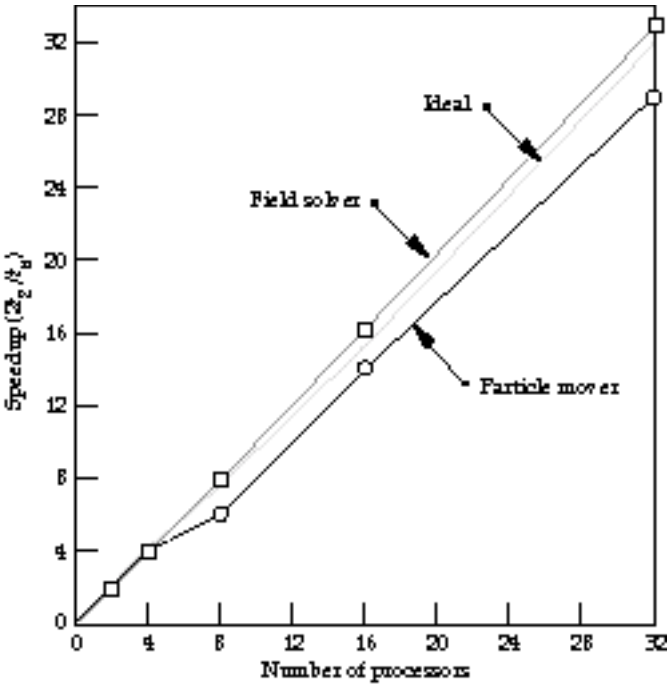

FIGURE 2. Plot showing the speedup of the code with an increasing number of processors. The speedup is twice the ratio of the time with two processors to the time with n processors, two being the fewest number of processors that the parallel code will run on. The light gray line shows the ideal linear speedup. The dark gray line with squares shows the speedup for the field solver. The black line with circles shows the speedup for the particle mover. Note that the speedup for the field solver is superlinear at first due to better use of cache memory with less data per processor. (08-10-0896-1907pb01)

Applications of WARP3d

WARP3d has been used extensively in the design and analysis of a number of experiments within the HIF program and in the study of the basic physics of the creation and transport of space-charge–dominated beams. In the following sections, we describe several of its most recent applications within HIF and an application outside of HIF.

Electrostatic Quadrupole Injector

A driver-scale injector and matching section is being built at Lawrence Berkeley Laboratory, in collaboration with LLNL, to address beam physics and technology issues.[15] The injector uses a sequence of electrostatic quadrupole (ESQ) lenses with a superposed voltage gradient along the axis. The net effect is to confine the beam transversely and to accelerate it. The beam is emitted from a hot-plate source and accelerated through a diode section into the quadrupole lenses. Figure 3 is an image of the injector when it is filled with the beam. Following the injector is a "matching section," a series of electric quadrupoles with decreasing aperture that reduce the transverse size of the beam down to "match" it into the following section of the accelerator.

WARP3d was used to understand a major issue in the injector design—the degradation of beam quality (brightness) by the nonlinear multipole components of the focusing fields and the so-called "energy effect." The energy effect arises from the spread transversely in the axial velocity of particles in the ESQ, which is due to the transverse variation of the electrostatic potential. Proper treatment of the beam dynamics requires detailed 3-D simulation with a realistic lattice model.

For design of the injector optics, the steady-state beam behavior was of primary interest. For efficiency, runs were made in a quasi-steady-state mode, taking several particle timesteps between each field solution and continuing until the system converged to a time-independent state. These 3-D simulations of the injector required less than 3 min on a single processor of the National Energy Research Scientific Computing Center's CRAY Y-MP C90 (at LLNL). This was fast enough to allow simulation of a large number of designs to optimize the injector, minimizing the beam degradation. Taking advantage of the ESQ's two-plane symmetry, the dimensions of the field grid were typically 40 × 40 × 348. Typically, 100,000 particles and 500 timesteps were used. Fully time-dependent simulations of the transient behavior require several hours of C90 CPU time.

The simulations were used as a guide in the development of a scaled, proof-of-principle experiment and then a full-scale experiment. In both cases, excellent agreement was found between the experiment and simulation results. Figure 4 shows the simulated injected beam in the full-scale diode. Figure 5 shows a comparison of the phase space representations of the simulated and the full-scale experimental beams.

For the matching section, the envelope equation solver in WARP3d is used to calculate the beam transverse rms size in the presence of quadrupole and uniform focusing multipole components extracted from the field solution, including the full geometry of the electric quadrupole conductors. Good agreement between the analytic envelope solution and the results of full PIC simulation implies that the envelope solver can be used for rapid iteration of the applied voltages to achieve the desired beam state. The full simulations are conducted to evaluate possible degradations in beam quality that cannot be evaluated in the analytic envelope model.

In a full-scale driver, multiple injectors will be used and their beams will be channeled into a multiple beam transport line by the matching section, resulting in a curved matching section. This system was simulated with WARP3d, combining three of its major capabilities: injection, a bent accelerator lattice, and field solution with complex conductor geometry. The simulations found no significant degradation of beam quality, thereby validating the design concept.

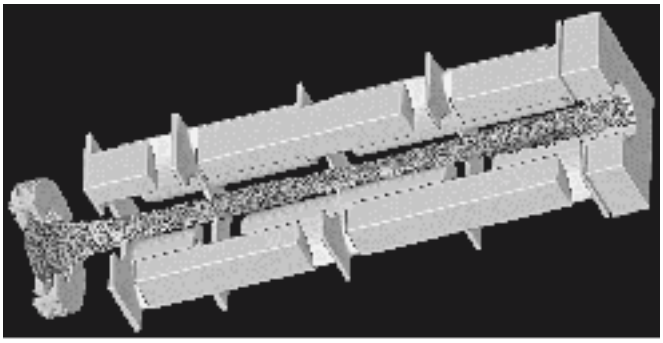

FIGURE 3. The 3-D image of the ESQ injector when filled with the beam. The shading of the beam particles is proportional to the particle energy relative to the energy at the beam center; the energy effect is evident in the darker shading of particles farther from the axis. (08-10-0896-1908pb01)

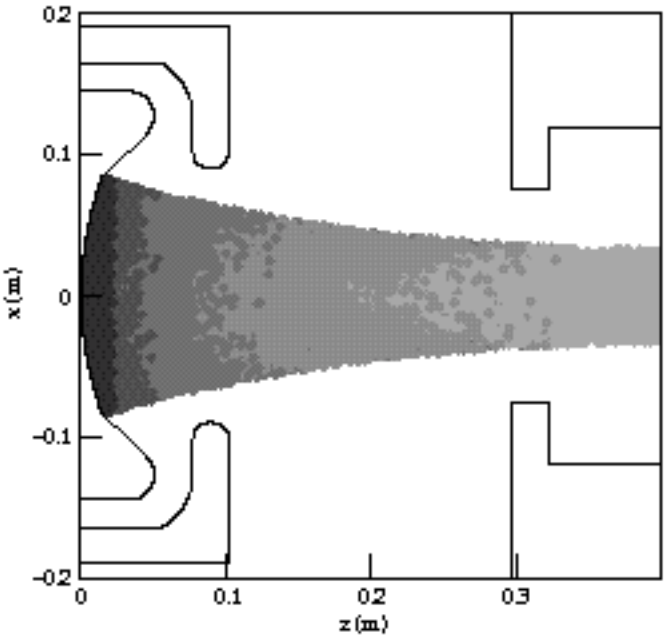

FIGURE 4. The diode of the ESQ injector filled with the beam. The outline is of the diode geometry. The shading of the beam particles is proportional to the time of injection of the particles, showing the darkest area injected last. (08-10-0896-1909pb01)

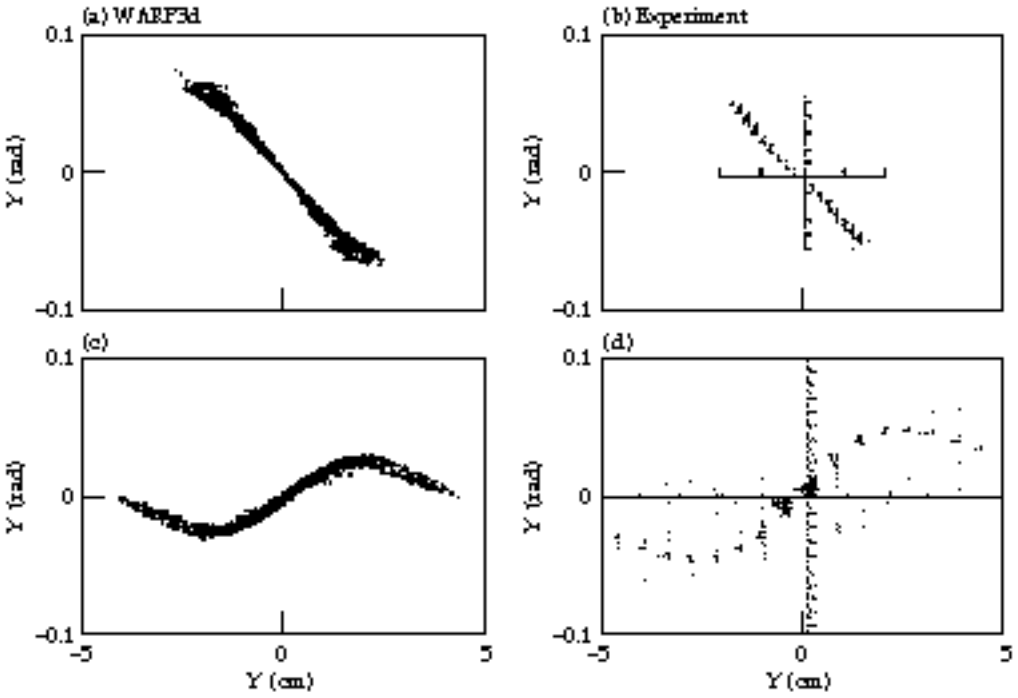

FIGURE 5. Comparison of experimental and simulation transverse phase space representations at the end of the ESQ injector. (a) and (b) compare the results at the design parameters. (c) and (d) compare the results with an increased diode voltage; in this nonoptimal case, the representations are significantly distorted from a straight line. (08-10-0896-1910pb01)

Recirculating Induction Accelerator

A recirculating induction accelerator, or "recirculator," offers the prospect of reduced cost relative to a conventional linear accelerator because the accelerating and focusing elements are reused many times per shot.[16] A small recirculator is being developed at LLNL to explore the beam dynamics of a full-sized fusion driver in a scaled manner; the key dimensionless parameters that characterize the beam are similar to those of a driver-scale ring, but the physical scale is much smaller.[17]

Although a full-scale driver will use magnetic dipole elements to bend the beam, electric dipole plates were chosen for the prototype to minimize costs. The size of the plates is constrained by space considerations and voltage-holding requirements, resulting in significant high-order fields. Using WARP3d, the plate shape was adjusted to minimize 3-D field nonlinearities and their influence on beam quality. Figure 6 compares the cross section of the beam after several laps with flat dipole plates and with shaped dipole plates.

WARP3d simulations were used to follow the beam for all 15 laps. Typically, the 2-μs-long beam is made up of 100,000 simulation particles. The field grid is a moving window covering four half-lattice periods. The grid measures 32 × 16 cells transversely (exploiting vertical symmetry) and 128 cells axially.

Figure 7 shows the time history of the normalized emittance at the midpulse for an acceleration schedule that does not compress the beam axially. The normalized emittance is a measure related to the transverse beam temperature and is given approximately by $\varepsilon_x \approx \tilde{x} v_x$

FIGURE 6. The beam cross section after several laps (a) with flat dipole plates and (b) with shaped dipole plates. The triangular shape in (a) is a result of the sextupole component from the flat dipole plates. (08-10-0896-1911pb01)

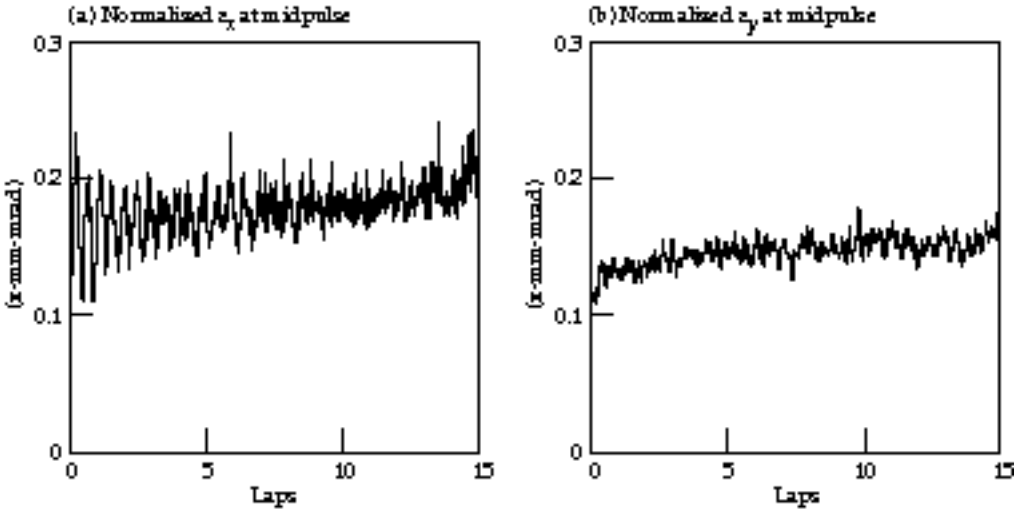

FIGURE 7. The time history of the normalized emittance and the beam midpulse show an acceptable increase. The initial oscillation and increase are due to axial thermal spread and the change from a straight to a bent lattice. (08-10-0896-1912pb01)

where x is a transverse dimension and $\tilde{x}$ is the rms of x. The simulations show an initial transverse heating of the beam on the order of 50% as the beam "matches" itself from the straight insertion line to the bent recirculator lattice. This heating is due to radial separation of particles with differing axial velocities in the bends, which leads to nonlinearities in the space-charge forces that thermalize to produce an increase in temperature. The observed growth agrees well with theoretical analysis.[18] After the initial heating, little further degradation in beam quality is seen over the 15 laps. The overall degradation is acceptable.

Near the end of the simulation, the emittance begins to rise due to the appearance of an instability that presents itself as a sinusoidal variation of the beam axial velocity along the beam. The variation begins to grow from the noise during laps 11 and 12 and becomes significant enough to affect the emittance during lap 15. The wavelength of the variation is on the order of the pipe radius. Numerous simulations seem to indicate that the instability is not numerical in origin; the wavelength and growth rate varied little with a wide range of numerical parameters, including timestep size and grid size. The instability requires all of the following: high space charge, bending, acceleration, and a uniform focusing force. The uniform focusing force results from the electric dipole plates. It is believed that the instability will not occur with magnetic dipoles, which will be used in a driver; the simulations made so far with magnetic dipoles have not shown the instability. The authors await the experimental results from the recirculator to see whether or not the instability indeed does occur as predicted.

Straight and Bending Experiment

As a precursor to the full recirculator experiment just described, the beam is being examined in a straight lattice and in a lattice with an 18° bend (using WARP3d). The straight experiment is being conducted to fully characterize the beam. The bent experiment will be the first detailed examination of a space-charge–dominated beam in a bent, alternating-gradient lattice.

The straight experiment consists of a series of electric quadrupoles and then magnetic quadrupoles. The field from the electric quadrupoles was calculated with WARP3d using the electrode geometry, and the field from the magnetic quadrupoles was calculated analytically. The quadrupole component was extracted from the fields and is used in the calculation of the envelope of the beam. Figure 8 shows a comparison of the

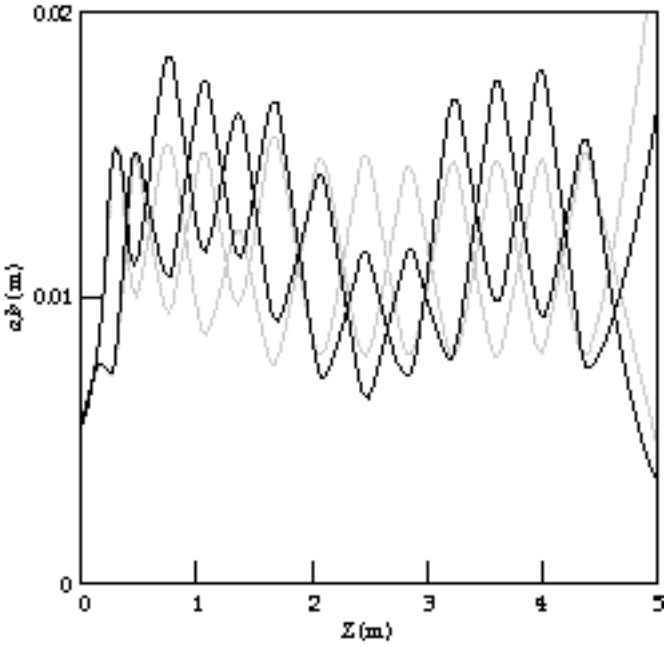

FIGURE 8. Comparison of the beam envelope as calculated with quadrupoles with hard-edged fields (gray) and with quadrupoles with calculated field profiles (black). The horizontal and vertical envelopes (*a,b*) for each case are shown in the same colors. The envelope was matched to the quadrupoles with hard-edged fields; this is seen in the nearly equal heights of the peaks in the gray curves. The significant difference shows the importance of including the correct quadrupole field profile. (08-10-0896-1913pb01)

envelope solution with hard-edge quadrupoles and with the calculated quadrupole data. Of course, given the detailed solution, it would be possible to construct a better hard-edged model.

The "simultaneous perturbation stochastic approximation"[19] optimization algorithm was implemented with the analytic envelope solver to find a matched envelope. The algorithm varies the voltages on the electric quadrupoles to find a matched beam in the first bent-lattice period. The envelope solver is being used by the experimentalists for the determination of the voltages for the experiment.

The bent beam experiment has been simulated with WARP3d to predict the amount of emittance growth from the bend correlated with the longitudinal thermal spread. Figure 9 shows the emittance growth for up to a 45° bend with a straight lattice afterward. Figure 10 shows the difference in the phase space before and after 18° using grayscales to show the relative longitudinal velocities of the particles. The simulations show that there should be an experimentally detectable amount of emittance growth in an 18° bend, assuming a longitudinal thermal spread of 0.05% of the axial velocity.

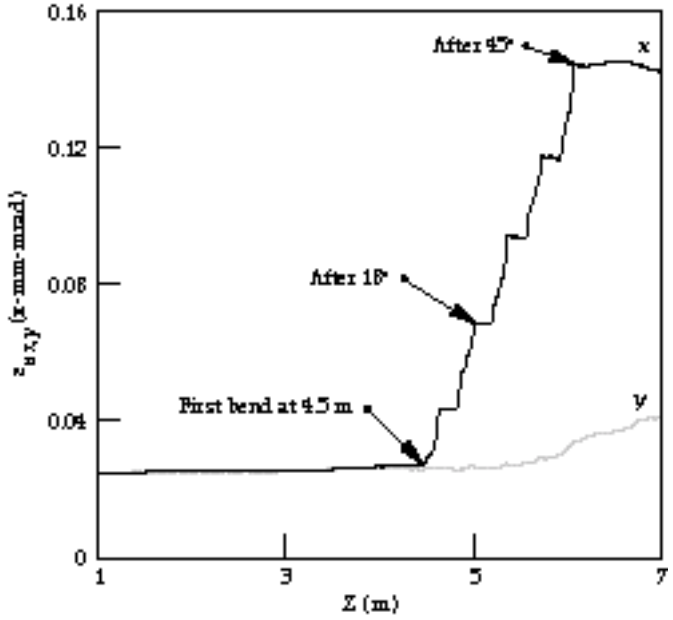

FIGURE 9. The WARP3d simulations show an emittance growth after 18° that should be experimentally detectable. (08-10-0896-1914pb01)

Sheet Beams

At the University of Wisconsin, WARP3d is being used in the examination of sheet electron beams formed by magnetic quadrupoles for use in high-power microwave sources.[20,21] The rapid axial variation of the beam profiles limits the use of the paraxial model and makes necessary the use of 3-D models. The 3-D envelope code TRACE-3D[22] is used for rapid iteration to design the system, but it does not model increases in emittance due to beam quality degradation. WARP3d simulations are used as a check of the envelope code design, to validate the envelope, and to examine the degradation of beam quality. An experiment is under construction and comparisons with simulation will be made.

Conclusions and Future Work

WARP3d was developed to examine high-current, space-charge–dominated beams for heavy-ion fusion and has proven its value in a broad range of applications, playing a critical role in the design and analysis of a number of experiments. Many features and capabilities have been implemented to allow inclusion of much of the relevant physics and simulation of large-scale problems.

Considerable further code development is planned, taking advantage of modern programming techniques such as object-oriented programming, and making more use of tools such as computer-aided design systems for geometry input, mesh generators, and massively parallel processing machines. These developments are geared to allow quicker and more complete analysis of ongoing experiments, leading to complete simulation of full-scale drivers.

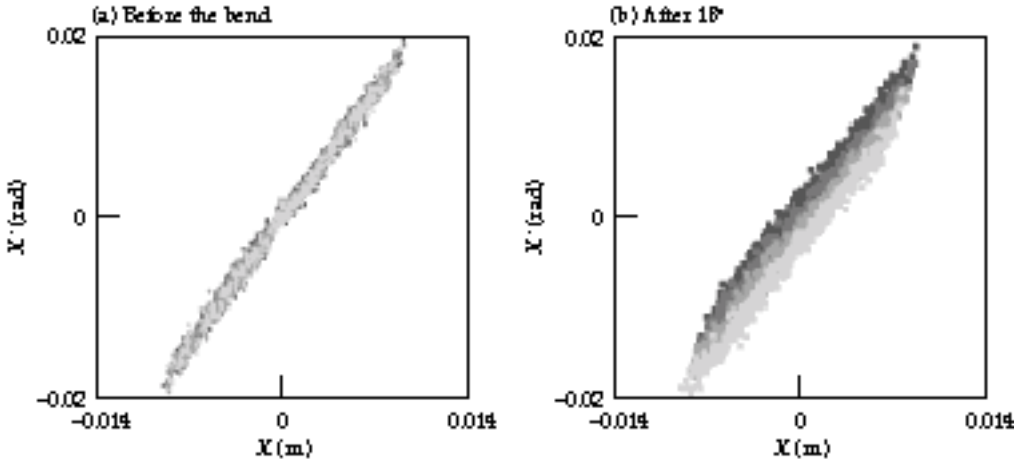

FIGURE 10. In the phase space, the transverse spreading of the beam due to axial thermal spread and a bend can be seen. The shading of the particles is proportional to axial velocity (lighter is faster). (08-10-0896-1915pb01)

Notes and References

1. R. O. Bangerter, *Il Nuovo Cimento* **106 A**(11), 1445 (1993).
2. A. Friedman, D. P. Grote, I. Haber, *ICF Quarterly Report* **1**(4), 115, Lawrence Livermore National Laboratory, Livermore, CA, UCRL-LR-105821-91-4 (1991).
3. A. Friedman, D. P. Grote, and I. Haber, *Physics of Fluids B* **4**(7), 2203 (July 1992).
4. William H. Press et al., *Numerical Recipes: The Art of Scientific Computing* (Cambridge University Press, 1986) pp. 652–659.
5. R. W. Hockney and J. W. Eastwood, *Computer Simulations Using Particles* (Adam–Hilger, Bristol, 1988) p. 215.
6. *TOSCA Reference Manual*, Technical Paper VF-11-92-14, Vector Fields Limited, 24 Bankside, Kidlington, Oxford OX5 1JE, England.
7. W. B. Herrmannsfeldt, *EGUN—An Electron Optics and Gun Design Program*, Technical Report 331, SLAC,1988.
8. K. R. Spangenburg, *Fundamentals of Electron Devices* (McGraw–Hill, New York, 1957).
9. A. Friedman and D. P. Grote, "Large-Timestep Techniques for Particle-in-Cell Simulation of Systems wth Applied Fields that Vary Rapidly in Space," *Proc. 1996 Comput. Accel. Phys. Conf.*, Williamsburg, VA, Sept 24–Oct 7, 1996.
10. A. Friedman, R. O. Bangerter, D. A. Callahan, D. P. Grote, et al., "A 3D Particle Simulation Code for Heavy Ion Fusion Accelerator Studies," in *Proc. of the 2nd European Particle Accelerator Conf.*, Nice, France (Editions Frontiers, Gif–sur–Yvette, France, 1990), p. 1699.
11. I. Haber, D. A. Callahan, A. Friedman, D. P. Grote, and A. B. Langdon, "Transverse-Longitudinal Energy Equilibration in a Long Uniform Beam," *Proc. of the 1995 Particle Accelerator Conf.*, Dallas, TX, pp. 3283–4 (IEEE, 1996).
12. I. Haber, D. A. Callahan, A. Friedman, D. P. Grote, and A. B. Langdon, "Transverse-Longitudinal Temperature Equilibration in a Long Uniform Beam,"to be published in *Proc. of the International Symposium on Heavy Ion Fusion*, Princeton, NJ, Sept 6–8, 1995, and in *Journal of Fusion Engineering Design*, 1996.
13. I. Haber, D. A. Callahan, C. M. Celata, W. M. Fawley et al., "PIC Simulation of Short Scale-Length Phenomena," *Space Charge Dominated Beams and Applications to High Brightness Beams*, S. Y. Lee, Ed., AIP Conf. Proc. No. 377 (AIP, New York, 1996) p. 244.
14. A. Giest et al., *PVM3 User's Guide and Reference Manual*, Technical Report ORNL/TM-12187.
15. S. Yu, "Heavy Ion Fusion Injector Program," *Proc. of the 1993 Particle Accelerator Conf.*, Washington, D.C., p. 703–705 (IEEE, 1993).
16. J. J Barnard, F. Deadrick, A. Friedman, D. P. Grote et al., *Physics of Fluids B*, **5**(7), 2698 (July 1993).
17. A. Friedman et al., "Progress Toward a Prototype Recirculating Induction Accelerator for Heavy-Ion Fusion," *Proc. of the 1995 Particle Accelerator Conf.*, Dallas, TX, pp. 828–830 (IEEE, 1996).
18. J. J. Barnard, H. D. Shay, S. S. Yu, A. Friedman, and D. P. Grote, "Emittance Growth in Heavy Ion Recirculators," *1992 Linear Accelerator Conf. Proc.* (Ottawa, Canada, AECL-10728, 1992) p. 229.
19. J. C. Spall, "Multivariate Stochastic Approximation Using a Simultaneous Perturbation Gradient Approximation," *IEEE Transactions of Automatic Control*, **37**, 332–341 (1992).
20. M. A. Basten, J. H. Booske, and J. Anderson, *IEEE Trans. Plasma Science* **22**(5), 960–966 (1994).
21. M. A. Basten, J. H. Booske, and S. M. Lund, "Three Dimensional Analysis of Electron Sheet Beam Formation Using Magnetic Quadrupoles," Lawrence Livermore National Laboratory, Livermore, CA; submitted to *Nucl. Instr. and Methods*.
22. *TRACE-3D Documentation*, Los Alamos Code Group, Los Alamos National Laboratory, Los Alamos, NM, Report LA-UR-90-4146 (1990).